U.S. DEPARTMENT OF ENERGY | Office of Science

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

Present and Future Computing Requirements for Polar Project & ACP SFA project.

Presenter: Jin-Ho Yoon
Pacific Northwest National Laboratory

NERSC BER Requirements for 2017
September 11-12, 2012
Rockville, MD

1. Project Description

- PI: Philip J. Rasch
- Key Players: Ghan, Easter, Fan, Fast, Ganguly, Liu, Ma, Ovchinnikov, Qian, Singh, Smith, H. Wang, M. Wang, Yoon, Zaveri

- ~ FY13: “Improving the Characterization of Clouds, Aerosols and the Cryosphere in Climate Models” --- ‘Polar project’ (multi-lab projects with LALN, LLNL, and LBNL).
  - Our NERSC activity is under this project currently.
- FY13 – FY15: “Interactions between Aerosols, Clouds and Precipitation in the Earth System” --- ‘ACP project’

Background: Aerosol cloud interactions and the Earth system.

Pacific Northwest
NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*

- *Aerosols are important and their sources are changing due to human activity*
- *Clouds are important to climate, they are strongly influenced by aerosols, and they strongly influence aerosol distributions*

- "Aerosol Indirect Effects" (AIE) are a major source of model uncertainty.
- Aerosols interact with many components of the climate system
(Clouds, Precipitation, Circulation features, Ice)

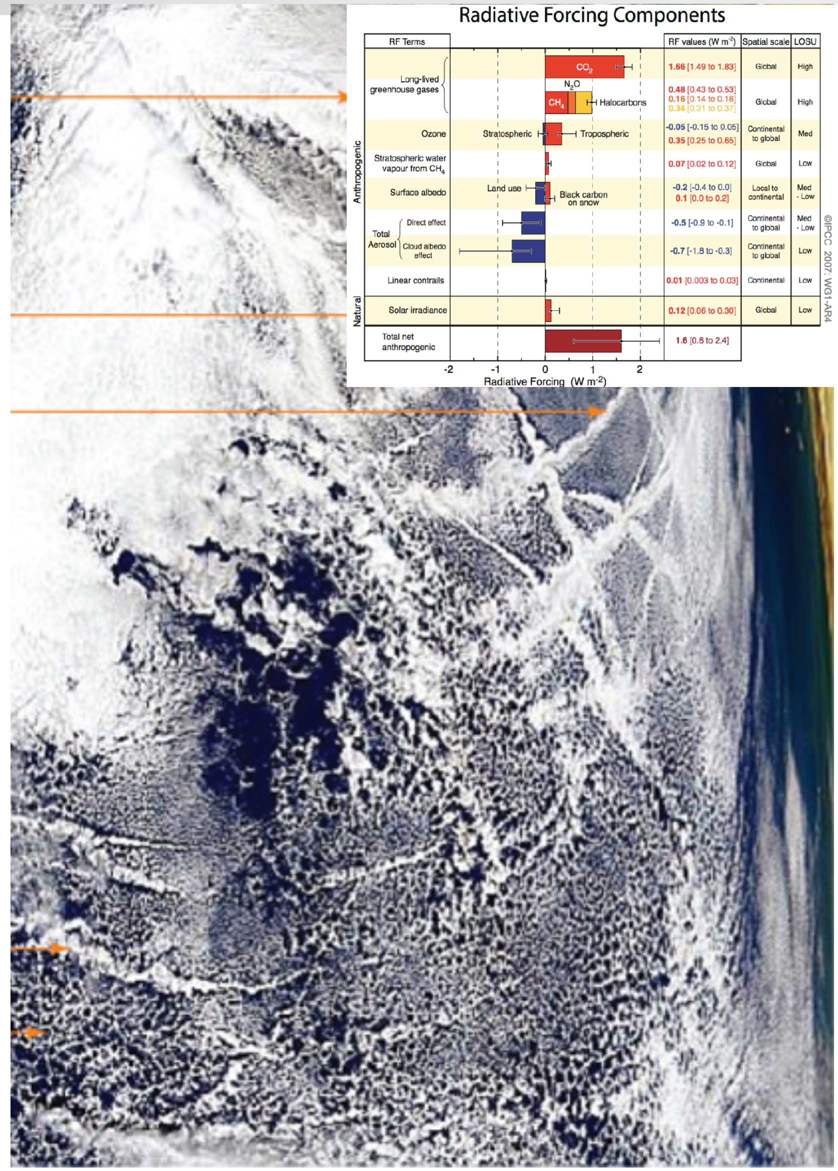

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

Scientific Objectives

Processes:

To improve and evaluate the methods for representing aerosols in climate models and quantify their direct and indirect impacts on the energy balance of the climate system, in particular through their influence on clouds, but also involving precipitation, including runoff, surface snow and ice.

Consequences:

To advance the understanding of how aerosols affect climate including the role of different components of the climate system in determining climate response to aerosol emissions.

2. Computational Strategies

- Our approaches are (1) to develop new physical schemes for climate model based on various observations, more physically based models such as PNNL-MMF, LES, and theory and (2) to use this climate models as a tool for further research.
- The codes we use are …
  - CESM1/CAM5: a major one
  - PNNL_MMF
  - LES (e.g., Wang et al. 2011)
- Our biggest computational challenges are…
  - Limit in scalability due to fv-CAM dycore → will be changed to CAM-SE soon.
    - Relatively small number of CPUs are used, but tend to go for more number of experiments.
  - Large amount of output due to long-simulations and more ensembles

2. Computational Strategies

- Our parallel scaling is limited by …
  - We have mainly used CESM1/CAM5 at 1.9 x 2.5 resolution (2deg).
  - Atmosphere only (CAM5 – 2deg): 960 Cores
  - Fully coupled (CESM1 – 2deg): 3000 Cores

- We expect our computational approach and/or codes to change (or not) by 2017 in this way …
  - From CAM5.2 (release on November 2012), CAM-SE will be default for 1-deg which has virtually no limit in scalability. However, Fv-CAM will be remained as a default at 2-deg.
  - Till now and for FY13, our focus is more on 2-deg resolution, but will be moving toward 1-deg eventually.

3. Current HPC Usage

- Most of the simulations are done on ‘hopper’.
  - Part of post processing is done on ‘euclid’.
- Hours used in 2012 (list different facilities)
  - In CY11, 17Mi (allocated) and 18Mi was used (including hopper’s free test).
  - In CY12, 14Mi (allocated as of now) and 10Mi has been used.
  - From Oct. 2012, we’ll use NCAR’s CISL new facility – Yellowstone
    - CSL proposal (PI: Phil Rasch) is granted for 15Mi hours
    - This one includes not only polar project (about 30%), but also EASM and others.
- Typical parallel concurrency and run time, number of runs per year
  - Atmosphere only (CAM5 – 2deg): ~20year/day
    - Usually short runs ~ 1 – 10 years.
  - Fully coupled model (CESM1 – 2deg): ~25 – 28year/day
    - Production runs for 20C runs (150years)

3. Current HPC Usage

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by **Battelle** *Since 1965*

- Data read/written per run
  - Most cases, we don't read much of input but write large amount of output.
  - We do have special configuration (data model experiments) that require large size of input, but it's been mostly tested on local HPC facilities at PNNL.
- CESM is well supported by NERSC staffs.
  - Any changes in software environments are notified through emails.
  - Website might be a better tool for communication.
- Data resources used (HPSS, NERSC Global File System, etc.) and amount of data stored
  - In CY11, 83,180 GBs stored in HPSS.
  - In CY12, 153,278 GBs stored in HPSS.
  - We are using 20TBs of project directory, but sometimes we have to utilize a couple of users' scratch directories all together.
    - Our project directory is not backed up due to size, but we do that regularly to HPSS.

4. HPC requirements for 2017

- Compute hours needed (in units of Hopper hours): Easily two to three folds.
  - If we use 1-deg (very likely) instead of 2-deg, it'll be easily more than four to six folds.
  - Our initial request for CY12 is 25Mi. By the end of 2012, we expect to get closer or exceed CY11 usage (18Mi).
  - In 2017, our request can be in the range of 50 – 60 Mi.
- Changes to parallel concurrency, run time, number of runs per year
  - So far, only a couple of users are heavily using NERSC facilities. But numbers of users will increase so the numbers of runs will increase as well.
- Changes to data read/written
  - As we go for higher resolution in space and time, our output will be not only two to three folds, but will be a lot more.
- Changes to necessary software, services or infrastructure
  - Faster I/O and larger disk space for postprocessing become more and more important.
  - Sharing data through ESG and more.

5. Strategies for New Architectures

- Our strategy for running on new many-core architectures (GPUs or MIC) is
  - Early tests of climate model show *little benefit* from GPUs
  - but there are ongoing efforts in a SciDAC project to improve this situation, and we are participants in that activity. We will take advantage of this capability if the project achieves success.
- To date we have prepared for many core b
  - Optimizing CPU configuration with help of P. Worley at ORNL.
- We are already planning to do …
  - CAM-SE & 1-deg
- To be successful on many-core systems we will need help with
  - Find out the best load balancing strategy. It is usually done at NCAR.
  - To get fine optimization on specific machine & configuration, support from NERSC is greatly helpful.

6. Summary

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

- What new science results might be afforded by improvements in NERSC computing hardware, software and services
  - Getting results faster and more ensemble members
  - Sharing data with others more efficiently
- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science
  - Faster I/O & larger disk space
  - A machine like 'euclid' for interactive jobs
    - Suppose we have 100TBs or PBs files, reading these and writing subset would take really long.

6. Summary

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle *Since 1965*

- NERSC generally refreshes systems to provide on average a 2X performance increase every year. What significant scientific progress could you achieve over the next 5 years with access to 32X your current NERSC allocation?
  - In translating climate information to societal decision making process, we do need to provide not just mean value, but also probability of events.
- General discussion